

NRC AI Regulatory Preparedness Introductory Remarks

Amy Sharp
U.S. Nuclear Regulatory Commission

ACRS Full Committee Information Briefing
September 3, 2026



Nuclear Energy Agency (NEA) RegLab

David Smeatham

Ennuvo

Matt Dennis

USNRC, Office of Research

International Regulatory Sandbox

ACRS Full Committee Meeting

**David Smeatham – Ennuvo Limited
Project Operating Agent**


**Faiza Sefta OECD – NEA – Project
Secretariat**

**Kevin Lee – CNSC – Project
Management Board Chair**

3rd September 2026



Innovation – why RegLab was developed in the first place



What is RegLab? A regulatory sandbox

Sandboxing is:

- Intended to reduce regulatory uncertainty.
- Create access to regulatory expertise.
- Help identify risk reduction mechanisms and approaches.
- Reduce the cost and time of entry for innovation.
- Allow regulators to test their arrangements.
- Test technology and process innovation in a controlled environment.
- Accelerate diverse and inclusive practices across the sector.
- Help develop capability.

Sandboxing is not:

- A short cut for existing regulatory processes.
- A means of changing regulation on a permanent basis.
- An endorsement from regulators for a specific technology or process.

In a RegLab, industry, regulators and academics work together to explore and test regulation of an innovation in a safe environment.

Table-top exercise.

Consideration of hypothetical, but representative, applications of innovation.



What value do we hope to achieve



A reduction in the regulatory uncertainty because we have regulators, licensees and technologies considering an innovation in a creative, safe environment.



Consideration of innovation at an international level to “bake-in” harmonised thinking at an early stage of the deployment of the innovation.



Help identify risk reduction mechanisms and identify where further work may be necessary so that the value of the innovation can be realised.



Allow regulators and licensees to test their arrangements.




Help develop capability.

International RegLab Project




The RegLab process



RegLab Pilot

- DIET Toronto – November 24, 2024
- Governance and human-in-the-loop aspects of two artificial intelligence applications
 - defining plant maintenance autonomously
 - generating safety documentation




EPRI – RegLab Pilot Report
(<https://www.epri.com/research/products/000000003002031833>)

International RegLab Project –RegLab # 1



- DIET Toronto – October 25, 2025
- AI to assist processes and procedures to enhance safety, security, safeguards and environmental protection, reducing costs of operations and maintenance.
 - AI substantiation / justification and explainability
 - Data assurance



Ennuvo – RegLab #1 Report
(<https://www.ennuvo.co.uk/s/RegLab-1-Final-Report-Mar-26-eno-040-V1.pdf>)

RegLab 1 Conclusions


- Participants agreed the use case could materially improve safety, security, safeguards and environmental protection (SSSE) and plant performance, while reducing operating and maintenance cost.
- Two dominant barriers to deployment were identified:
 - Explainability alone is not enough for high-consequence decisions. Explainability must be backed by defence-in-depth, including deterministic, physics-based models alongside the AI.
 - We need a shift in data mindset from “we have the data” to assured data quality, governance and representativeness to support a credible safety case.
- There are disbenefits if not done correctly:
 - Human complacency,
 - Skill degradation,
 - Introduction of new failure modes from added complexity, and
 - “Concerns about black-box” opacity.
- Five recommendations relating to (1) the need for AI good-practice guidance (including assurance framework, explainability, uncertainty, safety envelope, V&V, graded data quality); (2) close the AI skills/competency gap; (3) apply nuclear human-factors knowledge to AI-based decisions; (4) standardise data metadata/taxonomy for safe sharing; (5) evolve the RegLab process itself.
- The process worked: a logical, well-structured, inclusive “safe space” (Chatham House Rule) that enabled candid cross-stakeholder dialogue.

International RegLab Project –RegLab # 2




- NEA AI Conference – South Korea – May 26, 2026
- AI Tech Agents analyse plant data, compare it with licensing rules, automate condition-based tasks, and make competent-body-level technical judgements.
 - Autonomy – for the purposes of this use case, autonomy implies agents that operate independently, deciding which tools to use and how to achieve a goal after an initial, high-level prompt.
 - Scenario interrogates deployment of an approved, coordinated set of AI Agents and deterministic, logic-based, modules operating with independent management authority and responsible for the dynamic, end-to-end management of the INPO AP-928 Work Management cycle.
 - In the most autonomous scenario, the agent performs continuous monitoring and autonomous prioritization of the maintenance backlog. It independently executes the auto-release and scheduling of work without requiring real-time human intervention. The human role is only dedicated to governance and periodic auditing.

Supervised scenario



Autonomous scenario



RegLab 2 DRAFT conclusions

- Failure modes sit in the system, not just the AI model. The risk lives in the data feeding the AI and in the interaction between the AI, its operators and the wider plant.
- Human supervision depends on generic properties, not AI-specific ones: independent information, a structured decision process, contextual plant knowledge, and sustained competency. Left unmanaged, supervision decays into tick-box approval and the very competence it relies on erodes.
- Full autonomy does not shift SSSE duties away from the licensee. Rather, it changes how they are discharged, from real-time human judgement to up-front guardrails, verification & validation, and governance evidence designed in from the start.
- In general, the law is broadly sufficient. But while existing frameworks tend to cover AI used as a tool, the gap is mainly communication and capacity, not legislation. As AI becomes more agentic, the licensee centered laws may need to take account of responsibilities within the AI law's distributed, cross-border supply-chain obligations.

International RegLab Project –RegLab 3



- Virtual – Autumn / Fall 2026
- Considering remote operation of nuclear reactors to evaluate human factors, communication latency, and ensure robust safety oversight.

How can you get involved

- RegLab works by bringing a diverse group of people together to consider the deployment of innovation in a “safe space”
- We are always on the look out for people with relevant experience that they can share
- We are particularly keen to increase participation by non-represented parties (women, Indigenous group, etc. ..) and are seeking ideas on how to do so.





Thank you

David Smeatham
david.smeatham@ennuvo.co.uk

Faiza Sefta
faiza.sefta@oecd-nea.org

Kevin Lee
kevin.lee@cnscccsn.gc.ca

Near-Term Industry AI Use Cases

Matt Dennis, RES


Stephen Cumblidge, NRR

Norbert Carte, NRR

Josh Kaizer, NRR

Industry AI Use Cases



NRR Staff PoC	Related AI/ML Activity	Description
Stephen Cumblidge	Nondestructive Evaluation (NDE)	AI assistant to detect, measure, and classify weld defects and utilize virtual flaws for training
Stephen Cumblidge	Neutron Embrittlement (SRM-SECY-22-0019)	ML to develop neutron embrittlement trend curves for extended operational regimes capturing complex, multi-factor datasets
Norbert Carte	Online Monitoring (OLM)	OLM using AI to transition from time-based to condition-based instrument calibration technical specifications
Josh Kaizer	Critical Heat Flux (CHF) and other Special Computations	Machine learning regression techniques for predicting CHF and using AI/ML for data-drive alternative computational methods



<https://doi.org/10.1016/j.jnucmat.2025.156027>



<https://trueflaw.com/ml>



AI/ML for Non-Destructive Evaluation and Neutron Embrittlement

Stephen Cumblidge

USNRC, Office of Nuclear Reactor Regulation


AI/ML in NDE Data Analysis

- The variety of open-source ML tools has enabled the use of ML algorithms to be developed for many uses
- Work is ongoing by several groups to use these tools to develop automated data analysis algorithms to analyze NDE data
- These ML algorithms can potentially greatly increase the reliability and consistency of NDE in the field
- If used incorrectly, these tools could degrade inspections





Possible Benefits

- AI/ML analysis of NDE data has the potential to improve detection of flaws and improve the human factors of an examination.
 - In-service flaws are rare in the nuclear industry. Computers can maintain vigilance in cases where humans can be challenged.
 - Humans and computers make different types of mistakes, and a qualified analyst paired with an analysis run by ML gives the best of both worlds.
 - Reduced dose to inspectors if ML used to support manual UT examinations.



Proposed Implementation



- The approach suggested by EPRI, which is supported by the NRC staff, is for an AI/ML algorithm to flag areas with flaws, and the algorithm must find all flaws in the qualification set
 - The algorithm can produce more false calls than allowed in the given supplement
 - It will be up to the inspector to determine which of the areas flagged by the algorithm contain flaws, and ultimately the inspector is responsible for the results
- The NRC will continue to work with EPRI to monitor developments in the use of AI/ML for NDE data analysis
- Work described in “An Assessment of Machine Learning Applied to Ultrasonic Nondestructive Evaluation” (ORNL/SPR-2023/3245) ([ML24046A150](#))



From EPRI Report 3002023718

AI/ML for Neutron Embrittlement

- AI/ML tools can be very useful in analyzing complex data sets, such as for predicting neutron embrittlement in reactor pressure vessel steels
- Research, in cooperation with ORNL, will soon begin to:
 - Develop a description of the AI/ML data analysis of Neutron Embrittlement and a prioritized list and detailed description of the most important parameters in such an AI/ML algorithm
 - Develop performance-based metrics that can be used to verify and validate the results of AI/ML calculations of Neutron Embrittlement



AI/ML for Online Monitoring (OLM)

Norbert Carte

USNRC, Office of Nuclear Reactor Regulation

Charles McFeaters


USNRC, Office of Research

OLM based on ML

- OLM first discussed more than 20 years ago
 - In September 1998 EPRI submitted Topical Report TR-104965, “On-Line Monitoring of Instrument Channel Performance,” proposing the use of analytic redundancy and data-driven techniques to assess instrument channel performance and potentially extend calibration intervals
 - In July 2000 the staff issued a safety evaluation ([ML003734509](#)) with open items outlining 14 technical and programmatic requirements
 - NRC commissioned a series of reports focused on ML-based OLM ([NUREG/CR-6895](#) and [ML14274A306](#))
- In 2021 NRC endorsed an average-based OLM approach ([ML21179A062](#))
 - In 2022 NRC issues first LAR based on this approach ([ML23115A149](#))
- ML-based approach learns normal behaviour in context
- Flags abnormal behaviour for Engineering Evaluation
 - Both steady state and frequency response
- May require calibration

OLM based on ML

- TS surveillances ensure equipment behaves as expected
 - Calibration ensures accuracy
 - Monitoring ensures consistency of behavior
- Behavior between two calibrations defines normal
- Statistical uncertainty used to establish limit of normal
 - Can identify unacceptable drift or responsiveness



Example sensor calibration drift with respect to redundant peers ([ML24325A535](#))

AI/ML for Critical Heat Flux (CHF)

Josh Kaizer

USNRC, Office of Nuclear Reactor Regulation

AI/ML for CHF

- Across industry and research institutions, AI/ML techniques are being investigated to improve prediction of critical heat flux (CHF)
- In 2023, the OECD/NEA initiated a series of AI/ML CHF benchmark exercises ([ML25273A292](#)) as part of the Task Force on AI/ML for Scientific Computing in Nuclear Engineering ([link](#))
 - The exercise leveraged a large NRC-supplied dataset ([NUREG/KM-0011](#)) to test multiple ML approaches for improved CHF correlation accuracy
 - NRC staff contributed to the benchmark activity ([ML24207A085](#)) and are aware of commercial vendors that have expressed interest in using neural networks for CHF calculations



AI/ML for CHF

TRL 9	Proven in operation	✗
TRL 8	System complete and qualified	✗
TRL 7	Prototype demonstrated in operation	✗
TRL 6	Prototype demonstrated in relevant environment	?
TRL 5	Validated in relevant environment	?
TRL 4	Validated in laboratory	✓
TRL 3	Proof of concept demonstrated	✓
TRL 2	Technology concept formulated	✓
TRL 1	Basic principles observed	✓

- Staff have previously engaged in pre-application discussions about AI/ML methods for CHF and are currently awaiting the first application.
- Improved flexibility is proposed for changes to evaluation methods (10 CFR 50.221).
 - Proposed revisions to 10 CFR 50.59 and 53.1550 would allow licensees to implement certain changes to analytical methods without prior NRC approval, provided those changes are undertaken pursuant to an NRC-approved VVUQ program under 10 CFR 50.221, “Credibility requirements for modeling and simulation.” Proposed guidance is provided in DG-1468 (draft RG 1.271).
 - These proposals would shift NRC’s focus from reviewing each model or simulation directly, to assessing structured VVUQ programs determining model credibility.



AI and Autonomous Operation

Niav Hughes Green, PhD

USNRC, Office of Nuclear Regulatory Research

Matthew Anderson

Idaho National Laboratory

NRC Human Factors Research on Autonomous Operations

Niav Hughes Green, PhD

USNRC, Office of Nuclear Regulatory Research



Presentation Scope

- **Purpose:** Provide an overview of human factors research that supports regulatory preparedness for licensing reviews that include autonomous operations.
- **What this segment covers:**
 - Scope of current human factors research on **autonomy and autonomous operations**
 - The impact of higher levels of autonomy on **operator roles and responsibilities**
 - How human factors research supports consideration of **AI-supported autonomy**
- **What this segment does *not* cover:**
 - Part 57 or Part 53 requirements or policy decisions
 - Agency AI governance or AI use
 - Technical considerations outside of the human factors domain (e.g., instrumentation & control or cybersecurity)




Human Factors Considerations for Autonomous Operations

Research Initiated with Idaho National Laboratory to:

- Understand range of autonomous concepts anticipated
- Identify human factors considerations and implications
- Gather research findings and lessons from experience in other high-risk domains (e.g., oil and gas, maritime, aviation, and aerospace)
- Develop a taxonomy of human factors considerations important to ensuring safe operation
- Document technical basis and develop recommendations for guidance


HFE Guidance Development Process



Completed Reports



<https://www.nrc.gov/docs/ML2601/ML26013A039.pdf>



<https://www.nrc.gov/docs/ML2613/ML26139A285.pdf>



Key (Working) Definitions from RIL 2026-02

- **Autonomy:** The ability to operate independent of human control.
- **Automation:** A device or system that accomplishes (partially or fully) a function or task.
 - **Automated system:** A system that can operate independent of human control within fixed setpoints, parameters, or predefined scripts.
 - **Autonomous system:** A system that can operate independent of human control or intervention for extended periods in uncertain environments, with the ability to adapt to changing conditions, knowledge, and constraints. **An autonomous system may employ AI.**

Expanded (Working) Definition of Autonomous Operations

- *Autonomous operation* means the use of design features that enable a defined scope of operational and safety functions to be performed without reliance on human intervention, external command, or active control system input.
 - It may occur at the component, system, or plant level and may be used under normal, abnormal, or accident conditions.
 - It should not be equated to unmanned operation. Even highly autonomous concepts still assume human oversight and human accountability.
 - It could be achieved through inherent physics, passive safety features, deterministic automated logic, adaptive data-based algorithms (e.g., AI), or some combination of these.



Key Questions in Autonomous Operations

WHAT

- Is operating autonomously?
- What systems and/or functions?

WHEN

- Is it autonomous?
- Under normal, off-normal, emergency conditions?

HOW


- Is it autonomous?
- Active systems, passive systems, inherent characteristics?
- At different levels, functions, processes, or modes?

WHERE

- Is the autonomous operation controlled and/or monitored?
- Main control room, local control station, remote center?

Substitution Myth:

Automation simply replaces human action with automated action.





Human Factors Considerations with Automation



Situation
Awareness




Trust and
Automation
Reliability



Workload



Teaming



Training



Skill and
Knowledge
Retention



Novel Human
Error Types



Human Factors and Ergonomics Society Policy Statement

AI Guardrails for Safety Critical Applications

Explainable

Transparent

Tested with Humans

Safety Alerts

Fail Safe

User Training

Autonomy Validated

Design Considerations for Human-Centered Use of AI in Nuclear Power Operations




Explainable AI: Provide clear reasoning for AI outputs.



Human-in-the-Loop: Maintain operator authority for critical decisions.



Interface Design: Intuitive displays, alerts, and feedback mechanisms.



Training & Simulation: Prepare personnel for AI-assisted scenarios.



Concluding Thoughts

- Consider AI as a support tool, not a replacement for human judgment.
 - Clear guidelines for human-AI collaboration in nuclear contexts is needed.
 - Human factors evaluations should occur early and often during AI deployment.
 - Designs with AI should be explainable and transparent to the user.
- RES work positions the NRC to be ready for greater autonomy and potential use of AI.

DOE Remote and Autonomous Operation Research

Matthew Anderson

Idaho National Laboratory

3 Sep 2026


Idaho National Laboratory

Autonomous and Remote Operation Takeaways

Matthew Anderson, Jonah Lau, Michael Mendes, Vasileios Theos, Stylianos Chatzidakis, April Novak, Matthew Sgambati, Tim Grunloh

INL/MIS-26-95247

Reactors as geographically remote and autonomous distributed energy resource




Why?

- Remote operation may reduce costs (distributed controls of multiple reactors)
- Autonomous operations enable greater flexibility in power generation including load-follow
- Leverage momentum in advanced reactor deployment

Experimental Demonstration




Partial-Scale prototype
demonstration and
validation



AI Data Center Kinetics





- Hoodoo Lambda Hyperplane deep learning distributed memory system.
- 44 Nvidia A100 tensor core GPUs.
- Co-located autonomous (RL) and low autonomy (2-DOF) controllers.
- Power draw from nodes (aggregated to PUR-1 rating) on MNIST training and LLM inference served setpoint reference.




Reactor Kinetics

- 10 kW_{th}-rated PUR-1: Plate-Type MTR HALEU Fuel, Open Pool TRTR.
- Independent (from Reactor Safety System) I&C supports remote communication with Idaho-based controller.
- Co-located Digital Twin (DT) and cyber-physical testbed.
- Testbed's auxiliary moderator displacement rod (AMDR) induced power transients with up to 42 pcm of absolute reactivity.






- 1 Cyber-Physical Testbed
- 2 PUR-1 Rods
- 3 Testbed I&C System




- 1 Linear Actuator Support Structure
- 2 Linear Actuator
- 3 Linear Actuator Servo-Motor
- 4 PUR-1 Moving Platform




- 1 AMDR Guide Tube
- 2 Grid Location F7

Remote Autonomous operation of PUR-1




Seconds from start of experiment:
 11:21:06 AM (EST)
 9:21:06 AM (MDT)



Seconds from start of experiment:
 11:50:00 AM (EST)
 9:50:00 AM (MDT)

Autonomous Control Potentials

- Adaptive ML controllers coupled with DT state estimation
 - complex and emergent solutions within safe operational limits.
- Enable spatiotemporal anticipatory state-aware control.
- The “What’s” in state constraints:
 - Power-peaking factors, hot spots, structural degradation.
 - Accounting doppler and moderator effects, F-P poisoning, power distribution change.
 - Fuel burn-up and long-term state effects.




$$\hat{\mathbf{x}}^{DT} = \begin{bmatrix} \mathbf{x}^{\text{obs}} \\ \hat{\mathbf{x}}^{\text{lat}} \end{bmatrix}$$

$$\hat{\mathbf{x}}_{t+k}^{DT} \in \mathcal{X}_{\text{safe}}$$

Remote Operation Prospects

- Control loop frequencies evaluated at 1 Hz and 10 Hz (operational limit).
- Knowledge engineered data cataloging on cloud server supported centralized control.
- 1 Hz End-to-end round-trip latency within $\sim 2.5\%$ asymptotic period (~ 153 s) of AMDR with 5.24 ¢ reactivity.


[illegible]

Reactor Safeguards with Natural Language

- Control input and outputs are pseudo-anthropomorphically mapped.
- Agentic AI interface captures natural language command and embeds it into high-dimensional semantic-space with prior context.
- Facility's regulatory and operator intent context-aware c_t control is achieved:

π^* :

- $\mathbf{c}_t^{(1)}$: State-demand,
- $\mathbf{c}_t^{(2)}$: Anticipatory-operational,
- $\mathbf{c}_t^{(3)}$: Objective-constraint,
- $\hat{\mathbf{x}}_t^{DT}$: DT state estimate.




$\mathbf{u}_t \dashrightarrow \text{AMDR}$

$$\pi^* = \arg \max_{\pi \in \Pi} J(\pi; \hat{\mathbf{x}}_t^{DT}, \mathbf{c}_t), \quad (\hat{\mathbf{x}}_{t+k}^{DT}, \mathbf{u}_{t+k}) \in \mathcal{Z}_{\text{adm}}(\mathbf{c}_t), \forall k$$

AI Project Plan Update

Mock AI Tabletop Exercise

- Objective is to perform research and development of a comprehensive safety and security evaluation framework for the use of AI in NRC-regulated activities
- Tasks
 - Mock AI nuclear application safety analysis (Argonne National Laboratory)
 - Independent technical evaluator conducting a safety evaluation (Sandia National Laboratories)
- Research Information Letter (RIL 2026-02) publication September 2026
- Results used to further AI Regulatory Gaps Working Group deliberations



NRC AI Workshops and Symposia



U.S. Nuclear Regulatory Commission

2026 Artificial Intelligence Workshop

Tuesday, September 15, 2026

9:00 a.m. - 5:00 p.m. ET

Virtual or NRC Headquarters, Rockville, MD



<https://www.nrc.gov/public-involve/conference-symposia/data-science-ai-reg-workshops>

Moving Forward

- **Vigilance on AI:** Monitoring AI's integration in the nuclear sector and assessing regulatory needs
- **International Cooperation:** Leveraging global partnerships to accelerate safe, trusted, and innovative AI adoption
- **Proactive Engagement:** Identifying challenges, gaps, and use cases in AI technologies
- **Skill Development:** Ensuring staff are equipped to leverage and regulate emerging technologies

